

Advances in Time-Scale Algorithms

S. R. Stein
Timing Solutions Corporation
1025 Rosewood Avenue, Suite 200
Boulder, CO 80304

Introduction

The term clock is usually used to refer to a device that counts a nearly periodic signal. A group of clocks, called an ensemble, is often used for time keeping in mission critical applications that cannot tolerate loss of time due to the failure of a single clock. The time generated by the ensemble of clocks is called a time scale. The question arises how to combine the times of the individual clocks to form the time scale. One might naively be tempted to suggest the expedient of averaging the times of the individual clocks, but a simple thought experiment demonstrates the inadequacy of this approach. Suppose a time scale is composed of two noiseless clocks having equal and opposite frequencies. The mean time scale has zero frequency. However, if either clock fails, the time-scale frequency immediately changes to the frequency of the remaining clock. This performance is generally unacceptable and simple mean time scales are not used.

This paper will first review previous time-scale developments and then present some new methods that result in enhanced performance. The historical perspective is based upon several time scales: The AT1 and TA time scales of the National Institute of Standards and Technology (NIST), the A.1(MEAN) time scale of the US Naval Observatory (USNO), the TAI time scale of the Bureau International des Poids et Measures (BIPM), and the KAS-1 time scale of the Naval Research Laboratory (NRL). The new methods have been incorporated in the KAS-2¹ time scale recently developed by Timing Solutions Corporation. The goal of this paper is to present time-scale concepts in a nonmathematical form with as few equations as possible. Many other papers and texts discuss the details of the optimal estimation techniques that may be used to implement these concepts.

Clock Models

The perfect integrator is widely used as the mathematical model representing a precision clock. The fundamental model describes the continuous states of the clock. The time (or phase) is the integral of the frequency and the frequency is the integral of the frequency aging. For

¹KAS is a trademark of Timing Solutions Corporation

Report Documentation Page			<i>Form Approved OMB No. 0704-0188</i>		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 1992	2. REPORT TYPE	3. DATES COVERED 00-00-1992 to 00-00-1992			
4. TITLE AND SUBTITLE Advances in Time-Scale Algorithms		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Timing Solutions Corporation, 1025 Rosewood Avenue, Suite 200, Boulder, CO, 80304		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADA267301. 24th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, McLean, VA, 1-3 Dec 1992					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

some clocks, the frequency aging state may be zero for all time and can be deleted from the model. Most often, the continuous model is integrated to form a discrete model that is an explicit function of the time between observations. Let us call the phase-time of the clock x , the dimensionless frequency of the clock y , and the frequency aging, w . Then we write the discrete deterministic model as follows:

$$x_i(t+\delta) = x_i(t) + \delta y_i(t) + \frac{\delta^2}{2} w_i(t) \quad (1a)$$

$$y_i(t+\delta) = y_i(t) + \delta w_i(t) \quad (1b)$$

$$w_i(t+\delta) = w_i(t) \quad (1c)$$

We will also use two special symbols: $\hat{x}_i(t|t)$ is the estimate of clock i based on all data through time t and $\hat{x}_i(t+\delta|t)$ is the forecast of clock i at $t+\delta$ based on all data through t . The perfect integrator deterministic model is not always used. Several investigators have used the ARIMA modelling technique to study clocks [1,2]. This approach may use the perfect integrator model, or it may use model identification techniques to fit a model to the observations.

There are several methods of including the effects of noise on the system. One way is to add a random shock on the right hand side of each line in Eq. 1. These shocks are assumed to be uncorrelated in time and cause each clock state to behave as a random walk. This type of noise model is most commonly used for both Kalman filtering and ARIMA modelling of precision clocks. An alternative noise model was used as the basis for A.1(USNO,MEAN). In this model the frequency is given by a series of steps plus white noise. The frequency steps occur at random times. Mixtures of the two approaches have also been used and there is not yet a definitive answer to the question of what is the optimal noise model for precision clocks.

Review of Prior Techniques

With only one exception, all the major time-scale algorithms use the same approach to combine the individual clock times into the time scale. The current estimates of the clock times with respect to the time scale are determined by the requirement that the weighted sum of the differences between the current time estimates and their predicted values is zero. The mathematical expression of this requirement is called the basic time-scale equation:

$$\sum_{i=1}^N a_i \hat{x}_i(t+\delta|t+\delta) = \sum_{i=1}^N a_i \hat{x}_i(t+\delta|t) \quad (2)$$

If optimal estimation techniques are used as the basis for prediction, the differences between the current estimates and the predictions form a white sequence and averaging is the appropriate method of combining the contributions from different clocks. Combining the residuals instead of the clock times themselves results in time continuity upon addition or deletion of clocks. The basic time-scale equation is also responsible for establishing the weights, a_i , that control the contribution of each clock to the time scale.

The only time-scale algorithm that doesn't use the basic time-scale equation is TA(NIST) and a similar algorithm called the Composite Clock. This algorithm defines state equations

that describe the time evolution of the states of the ensemble members including the effects of noise. Optimal statistical techniques are used to estimate the times of each constituent clock. However, the results are not good. In fact, the estimation problem of TA(NIST) is underdetermined. This situation results from the failure of the algorithm to implement the basic time-scale equation or some alternative and is not the result of the statistical estimation techniques (Kalman filter) that are employed. As a result, the effective weighting that controls how the clocks are combined is different from the intuitive weighting schemes used in the other algorithms and can't be controlled by the user. In addition, the time scale is hard to manage. Addition of new clocks is a particular problem. Another time scale called the Composite Clock is used to compute GPS time. It is based on the TA(NIST) approach and suffers from the same problems.

The question of how to form the forecasts needed by the basic time-scale equation has been ignored throughout the preceding discussion. Unfortunately, the same question was not dealt with sufficiently carefully by any of the time-scale algorithms. Instead, ad hoc procedures have been used to estimate the clock frequencies used in predicting the future clock times. The ad hoc procedures limit time-scale performance. In particular, these algorithms do not solve the problem of how to form a time scale that has optimum performance for both long and short sampling times when some clocks have good short-term stability and others have good long-term stability. Recent developments, described in the New Techniques section, have eliminated this limitation.

New Techniques

The problem of determining the clock times and frequencies would be trivial if it were possible to make absolute time measurements of the clocks. However, it is only possible to measure the time differences between pairs of clocks. This situation may be described by saying that the states of the clocks are not observables of the system. One way to view the problem is as follows. Each clock is perturbed by two noise inputs at every observation. One random shock contributes to the evolution of the clock time and the other to the evolution of the clock frequency. Suppose that there are N clocks that have no noise on the frequency. Then there are $N-1$ clock time difference measurements. One cannot uniquely determine the N random shocks that perturb the time. However, when one uses the basic time-scale equation in addition to the $N-1$ measurements, there are N relations among N unknowns and one obtains a unique estimate for each random shock and therefore for the clock times themselves.

At first it might appear that, knowing the times, one may unambiguously calculate estimates of the clock frequencies. However, this is not the case when there are independent noises on both the time and the frequency states. The time difference measurements, together with the basic time-scale equation, suffice to determine the total noise input to each clock's time. However, that total noise input is composed of a direct contribution to the time and a contribution via the noisy frequency. There is insufficient measurement information to separate these two noise sources. Thus, with certain limitations, there is freedom in determining the clock frequencies. For example, one could require that all the frequencies are zero for all time and still be consistent with all the measurements. The problem is how to allocate the

noise between the time and the frequency states.

The allocation of the noise between time and frequency is partially accomplished by requiring that the model of each clock has the appropriate statistics. This requirement is sufficient, for example, to exclude the case of all clocks having constant frequency but is not sufficient to allocate the frequency noise on a detailed basis in such a way that the time scale has all the desired properties. If nothing further were done, the frequency noise would be allocated according to the weights used in the basic time-scale equation. Independent detailed allocation of the frequency noise is accomplished by adding a supplemental time-scale equation that also guarantees that the time-scale frequency is continuous when clocks are added to or dropped from the time scale. By analogy with the clock times, one requires that the weighted sum of the differences between the current frequency estimates and their predicted values is zero. The same technique must be extended to the frequency aging state if it is nonzero. The supplemental time-scale equations are written:

$$\sum_{i=1}^N b_i \hat{y}_i(t + \delta | t + \delta) = \sum_{i=1}^N b_i \hat{y}_i(t + \delta | t) \quad (3)$$

$$\sum_{i=1}^N c_i \hat{w}_i(t + \delta | t + \delta) = \sum_{i=1}^N c_i \hat{w}_i(t + \delta | t) \quad (4)$$

The basic time-scale equation and the supplemental time-scale equations have been incorporated in the KAS-2 time-scale algorithm developed by Timing Solutions Corporation. Using this approach, each clock's contribution to the time-scale is determined by a set of weights, one for each state that is perturbed by noise. There are always at least two weights per clock. As a result KAS-2 may be independently optimized in both the short term and the long term. It can therefore utilize clocks with very different frequency stability without degrading the stability of the time scale. U. S. Patent 5,155,695 has been issued covering time scale systems utilizing the supplemental time-scale equations. Other patent applications are pending.

Implementation

There are two basic approaches to time-scale algorithms — real time *v.s.* after the fact. Real-time algorithms are required for applications that need a physical signal derived from the time-scale. This is true, for example, if the time-scale is used to provide stable and fault tolerant time and frequency for a communications network. Other applications are amenable to post processing of the data. In this case it is possible, in principle, to compute time-scale estimates with smaller mean-squared errors since more data are available. TA(NIST), A.1(MEAN,USNO), and ALGOS(BIPM) are computed after the fact. They do not use smoothing techniques but they do use both subjective and objective techniques to detect "unhealthy" clocks and modify the ensemble membership before performing the final time-scale estimates. KAS-1(NRL), KAS-2(TSC), and AT1(NIST) all produce estimates in real time. Other experimental algorithms utilize smoothing techniques to attempt to improve their estimates compared to a real-time filter.

Optimal State Estimation

The deterministic clock model, the noise model, the basic and supplemental time-scale equations, and the observations are sufficient to calculate the time scale. Since the problem is stochastic in nature, the solution is not unique. However, there are several methods for calculating minimum squared-error estimates of the clock states. Among these are: least squares estimation, Bayesian estimation, maximum likelihood estimation, Kalman filtering, and Wiener filtering. All these techniques are equally good in the sense that they minimize the mean square estimation error as long as the assumptions are the same in each case [3].

Robust Statistics

Inevitably, the data presented to the time-scale algorithm will be deviant in some way. Perhaps one of the clocks changes characteristics dramatically due to an internal failure or perhaps the observations are corrupted during a data transfer. It is desirable for the time scale to be robust under these conditions. By this we mean that deviant behavior in a small number of observations does not unduly influence the performance of the time scale [4].

Robustness is ensured by a two step process. First we generate an estimate of the expected observations and then we act to reduce the effect of any observation that is far from our expectation. Both parts of the process must be robust. Consider the problem of detecting outliers in a set of measurements of a constant such as the length of a rod. One might think to compare each measurement with the mean of the set. However, when one measurement is bad, the mean may differ arbitrarily from its expectation value making the deviation from the mean an undesirable measure of deviant behavior. The median is much more robust.

The most common method of dealing with outliers is to reject observations that appear deviant. But this procedure leads to a discontinuity in the estimation procedure where small changes in an observation can produce significant changes in the resulting estimate. This discontinuity can produce transients and even instability and should be avoided. Continuity is ensured by continuously deweighting observations over a range of values.

Finally, it is sometimes argued that if one has a wealth of data, it is desirable to discard large amounts of good data to enhance the probability of rejecting the bad data. This must not be done if the residuals from the estimation process are used to determine the noise of the clocks. If good data are discarded under these circumstances, the clock noise will inevitably be underestimated.

Parameter Estimation

The function of time-scale algorithms is to allocate the observed clock noise among the contributing clocks. This can only be done if the parameters of the clock models of each clock are accurately known. These parameters are the spectral densities of each of the noise sources that perturb the clock and any of the commonly used techniques for estimating the noise of clocks may be applied. For example, one may characterize the Allan variance of clock pairs and use the three corner hat technique to separate the variances. The principal drawback of this approach is that Allan Variance estimation requires uniformly sampled data.

The Maximum Likelihood method [5] is a more sophisticated approach. It relies on the fact that the likelihood function computed from the residuals of the time-scale estimation process is maximized when the correct parameters are used. The model itself can also be checked for validity by this method since the correct model used with the correct parameters results in Gaussian white residuals. Standard statistical techniques can be used to analyze the residuals for whiteness. The principal drawbacks of the maximum likelihood method are that it requires storing large amounts of historical data and substantial processing time.

Other techniques have been developed for real-time algorithms that are incompatible with storing and processing large amounts of historical data. For example, AT1(NIST) uses the one step ahead prediction error as its sole measure of clock noise. The value is updated recursively in an exponential filter after each cycle of time-scale computation. An extension of this technique is the analysis of the one step ahead prediction error as a function of the prediction interval [6]. Since different noise sources dominate the time prediction error as the interval varies, it is possible to estimate all the important noise spectral densities of each clock.

Performance of the KAS-2 Time-Scale Algorithm

The best way to evaluate the performance of a time-scale algorithm is to simulate the clocks in the ensemble since simulation makes it possible to know the true times of the clocks. There are several appropriate simulation techniques [6,7]. The method developed by this author has been used to test KAS-2. However, simulation testing should be supplemented with tests using actual clock data. The comparison of several different time scale algorithms on real clock data provides the ultimate reassurance that something important has not been overlooked.

Figure 1 shows the simulated performance of KAS-2 with an eight clock ensemble composed of four clocks with good short-term stability and poor long-term stability and four additional clocks with poor short-term stability and good long-term stability. The resulting time scale is better than the best clocks everywhere.

Figure 2 compares the times of three independent time scales computed from the measurements of the clock ensemble at the National Institute of Standards and Technology using data supplied by NIST. The test period was the first 10 months of 1992. The AT1(NIST) time scale was computed in real time while TA(NIST) was computed after the fact using the information generated by AT1 to adjust the ensemble. KAS-2(TSC) was run on the historical data as if it were running in real time. It used the same estimates for the clock noises that were used in TA(NIST). All three scales were adjusted so that they had the same rate as TAI at the beginning of the year. In addition, an intentional frequency steer of 3.98 ns/day in AT1(NIST) was removed in order to compare the free time scales. One should not jump to conclusions since there is only one year of comparative data, but the results certainly support the theoretical claims of the advantages resulting from the implementation of the supplemental time scale equations.

Figure 3 illustrates the performance of a Cs clock at the Naval Research Laboratory that was steered in real time to the KAS-2 time scale. The time scale was computed every 5 minutes and approximately one-tenth of the estimated phase error was removed each sample

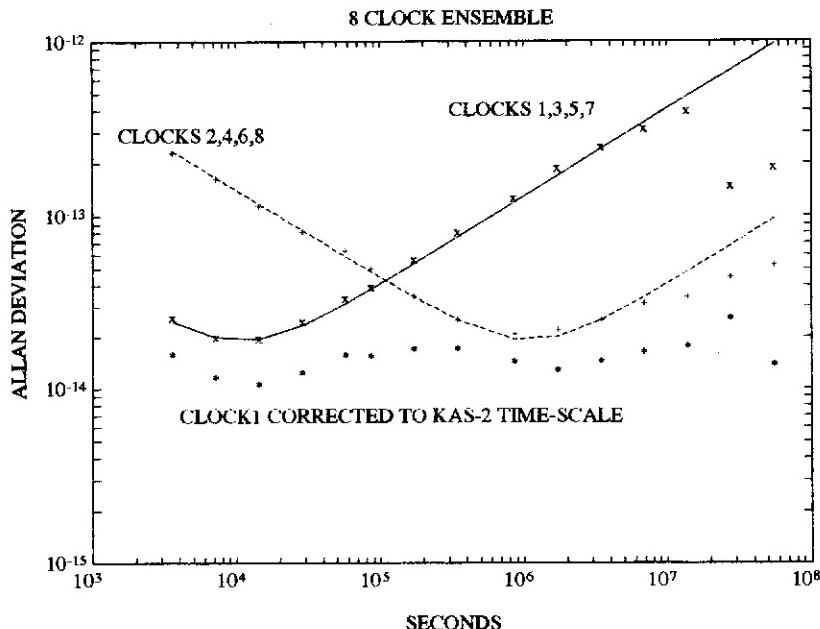


Figure 1: The simulated frequency stability of a time scale composed of four clocks with good short-term stability and four clocks with good long-term stability.

by a second order phase-lock loop. As a result, the frequency stability of the Cs standard equals its free running value for times shorter than three thousand seconds (10 sample intervals). At longer times, the frequency stability improves until it reaches the performance of the KAS-2 time scale.

Comparison of Time-Scale Algorithms

Table 1 summarizes the features of the time-scale algorithms discussed in this paper. Some of the more important differences and similarities will be discussed in more detail here.

All the algorithms except for TA(NIST) and the similar GPS Composite Clock implement the basic time-scale equation. Neither do these two algorithms implement a suitable alternative. As a result, they both perform worse than the others. First of all, they do not provide user control of the weights that determine the clock contributions to the time scale. Second, they have difficulties when clocks enter or leave the ensemble. Finally, they are likely to be unstable when adaptive filtering is used.

The algorithms that implement the basic time-scale equation fall into two categories. TAI(BIPM) and A.1(USNO,MEAN) are computed after the fact. In practice, human judgement supported by objective statistical analysis is used to filter the input data in an attempt to improve the time-scale performance. AT1(NIST), KAS-1(NRL), and KAS-2(TSC) are real-time time scales. In principle, after the fact time scales are capable of performing bet-

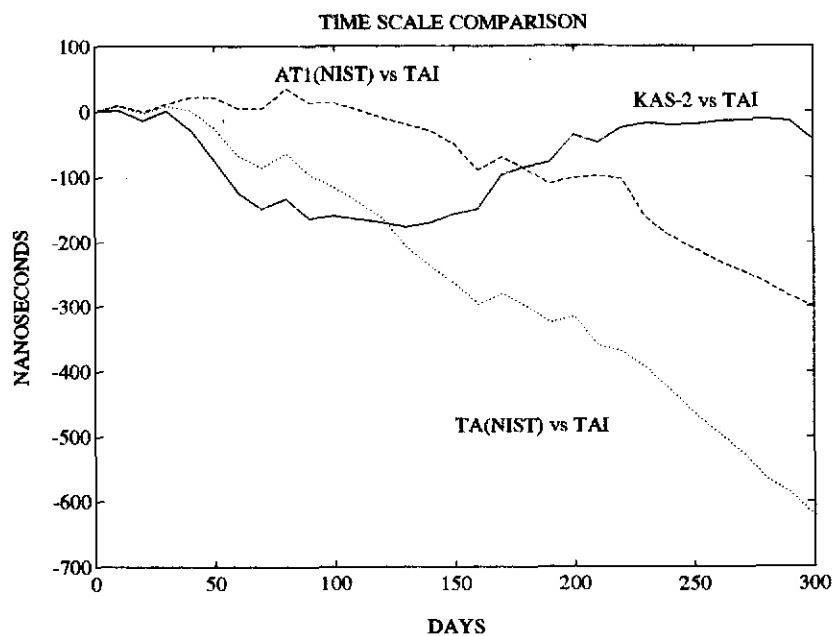


Figure 2: The times of KAS-2, AT1(NIST), and TA(NIST) *v.s.* International Atomic Time (TAI) for the first 10 months of 1992.

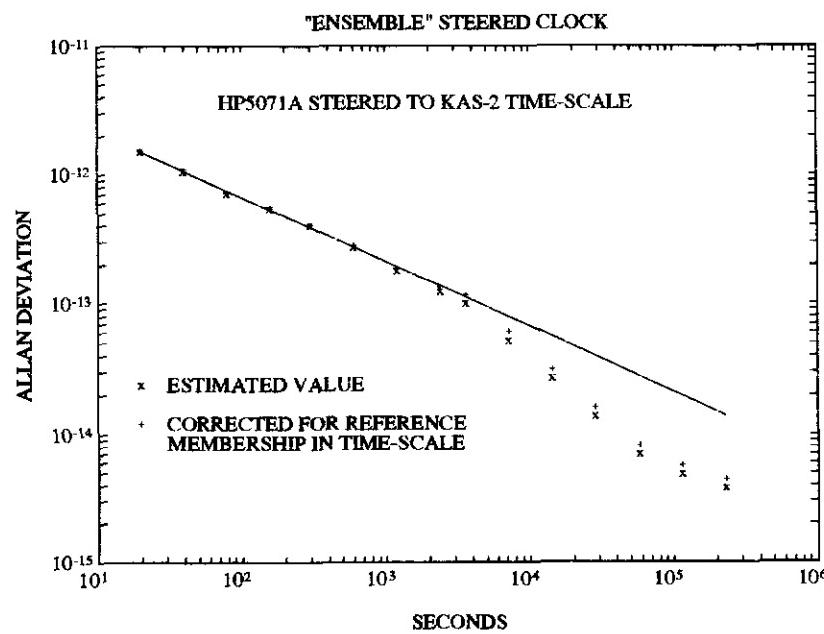


Figure 3: Stability improvement of a cesium clock steered to the KAS-2 time scale.

Table 1: Comparison of time-scale algorithms

	USNO A.1	BIPM ALGOS	NRL KAS-1	TSC KAS-2	NIST AT1	NIST TA
deterministic clock model is an integrator	✓	✓	✓	✓	✓	✓
uses explicit noise model	✓		✓	✓		✓
uses basic time-scale equation for the time state	✓	✓	✓	✓	✓	
uses supplemental time-scale equation for the frequency state					✓	
uses supplemental time-scale equation for the frequency aging					✓	
protects against outliers	✓	✓	✓	✓	✓	✓
uses robust statistics			✓	✓		
accounts for measurement noise			✓	✓		✓
estimates true clock performance				✓		✓
estimates weighted time-scale		✓	✓	✓	✓	
uses separate weights to optimize short, medium, long-term stability				✓		
accepts frequency measurements				✓		
solves an observable problem	✓	✓	✓	✓	✓	
accepts aperiodic data				✓		✓
time-scale available in real-time			✓	✓	✓	
available commercially				✓		

ter than their real-time counterparts because they have available more information. However, real-time time scales are necessary whenever the time scale is used to control physical clocks or for synchronization. Each approach has its advantages and should be used in the appropriate circumstances.

Many of these time scales differ in their implementation of robust statistics. The after the fact scales typically use rejection rule techniques to exclude bad clocks entirely. Use of rejection rules in real-time algorithms is inappropriate since the discontinuity of the rejection rule introduces a potential instability that can result in rejection of all the clocks in the ensemble and require human intervention to restart the time scale. AT1(NIST) has suffered from this problem. KAS-1(NRL) and KAS-2(TSC) eliminate this discontinuity by implementing Percival's outlier deweighting method. Large outliers are still rejected, but there is a continuous transition between full acceptance and full rejection of data.

Different clock models are used in the time scales discussed in this paper. The USNO A.1

time scale models the frequency noise as a series of discrete steps of random amplitude and time of occurrence. The other time scales model the frequency noise as either a continuous or a discrete random walk process. An advantage of the random walk model is the existence of straightforward filtering and smoothing techniques that can provide objective estimates that are optimum in the minimum squared-error sense. Several state estimation techniques have been used in these time scales. However, regardless of the state estimation method, the update equation used for the states is always an exponential filter. Thus, the only issue is the proper selection of the filter time constant. Despite the simple description of this problem, the solution is quite difficult except in special situations such as an ensemble consisting of several identical clocks. TAI(BIPM) and AT1(NIST) utilize restricted operating conditions and approximate formulas for the required filter lengths. The approach works well as long as the restrictions are not violated. TA(NIST), KAS-1(NRL), and KAS-2(TSC) utilize the Kalman filter approach. This guarantees that the filter lengths are correctly determined as long as the clocks fit the models and the parameters are correctly estimated. However, TA(NIST) and the GPS Composite Clock are often intentionally mismanaged by changing the parameters of the clocks since they do not have explicit user controlled weights.

Acknowledgements

The author wishes to thank the Naval Research Laboratory for prior support. Several people made large contributions to the testing and evaluation of the KAS-2 algorithm. G. A. Gifford at the Naval Research Laboratory tested the UNIX version and J. Levine at the National Institute of Standards and Technology tested the VAX VMS version. Their feedback was invaluable.

References

- [1] G. E. P. Box and G. M. Jenkins. *Time Series Analysis: Forecasting and Control, revised ed.* Holden-Day, San Francisco, 1976.
- [2] Donald B. Percival. The U. S. Naval Observatory clock time scales. *IEEE Trans. Instr. and Meas.*, IM-27:376–385, December 1978.
- [3] Arthur Gelb, editor. *Applied Optimal Estimation*. The M. I. T. Press, Cambridge, 1974.
- [4] Donald B. Percival. Use of robust statistical techniques in time scale formation. *Second Symposium on Atomic Time Scale Algorithms*, June 1982.
- [5] Peter V. Tryon and Richard H. Jones. Estimation of parameters in models for cesium beam atomic clocks. *NBS Journal of Research*, 80:3–16, January / February 1983.
- [6] L. A. Breakiron, G. A. Gifford, and S. R. Stein. Report on the timescale algorithm testbed at USNO. *Proc. 21st Annual PTTI Applications and Planning Meeting*, 269–287, December 1989.

QUESTIONS AND ANSWERS

Question: M. J. Van Melle, Rockwell International: I was interested in the table you had on comparison of time scale algorithms and also your CAS1 which was on your last chart. I am also working at (??) and you know the Air Force has partitioned all the GPS's into different partitions or ensembles. Up to six they can handle and now down to five and they have three partitions. But they can only work if they have the Cesium clocks in there. Once they put a Rubidium in there, it doesn't work. I was wondering if you could explain why that is the case.

Question: S. Stein: I really did not pay you enough for that question. The composite clock algorithm that IBM did is essentially the same as the NIST TA algorithm. It has some differences in the way they treat the renormalization of the covariance matrix to get around this problem of the covariance matrix blowing up, but fundamentally the problem with that type of algorithm is it implements the clock model — perfect integrator clock model — and it estimates the time of the clock based on the measurements, but it does not implement the basic time equation. The solution is not constrained to be the weighted average of the times of the clocks based on user selectable weights. That's the fundamental reason for what I would call instabilities you will see in the state predictions of algorithms of that type. So I think is what happens is you have both additional burden operationally on the Air Force personnel at the Master Control Station and try to maintain the algorithm and you have limitations on not being able to include different types of clocks with different characteristics. The algorithm only works on a limited set of circumstances. I do not know what to do about that except to fix that algorithm so indeed the basic time scale equation is incorporated.

Question: Richard Keating (USNO): Have you ever studied the sensitivity of the basic time scale equation as you call it to round off and loss of significance. Some remarkably disparate numbers in the sense of very large numbers and very small numbers play a role when you start averaging. I was wondering if you ever made a study of that or if you know anybody who has.

Answer: S. Stein: I have done to date approximately two hundred seventy simulations studies of the performance of these algorithms. The problem you raise is possibly the most difficult problem of all in the computation of system time. I am not sure I can remember the exact figures but the asymmetry in the covariance matrix occurs depending on the number of clocks and the difference in the performance of the clocks in the eighth or tenth digit while computing in double precision. Once one cycle through the recursion costs you eight digits of precision. Yes, this is a terrible problem because the fundamental approach to generating the gains, the filter time constants, is differencing large numbers against small numbers. You can treat that problem, for instance one of the ways of treating that problem and Kalman Filter methods is to go ahead and resymmetrize the covariance matrix each time. You also must treat the problem in special cases. For example we provide the facility for the user to carry along clocks which we call not members of the ensemble. That means that the clock is defined as not effecting the time scale computation, no matter how large a deviation it makes. It can move one second and shall not move the ensemble one nanosecond. There is simply not enough computational accuracy to do that, so what we do is we go in and we fix the Kalman gains directly to cope with that situation. You always have to do some special things, there are some special things we do, and that is one of them. Fortunately the statistical analysis that has been done for much more important applications than this for navigating spacecraft and flying airplanes, have dealt with all of these problems and you can use a more sophisticated version of a standard Kalman Filter that allows you to go in and change the gains and still correctly compute the recursion, and to correctly take into account the ad hoc changes you

have made in the gains. We do that.

Manuel Aparicio (ITT): I have two questions. The first one is if you have a non-evenly spaced sample for the Kalman Filter that you are using, you mention that you will have some problems. What type of problems are they? The second one is if your estimates of the noise processes in the Kalman Filter are not absolutely correct for one of the members of the ensemble, what kind of effects will you see?

S. Stein: The first is a potential problem. Box Jenkins or ARIMA modeling requires evenly spaced data. The typical way around that is if you are missing data you interpolate it. It biases your results. We use the approach that was done by Dick Jones and (Petrian?) for the TA algorithm and that is we use a noise model which is the integration of continuous white noise and so the Kalman Filter runs optimally for any spacing of data, therefore you do not need equally spaced data. Your second question. The effect of inaccuracies in the estimates of the parameters. Of course you never know the parameters – you only estimate them. The values of the spectral densities of clocks determine the filter time constants. The bottom line, forget everything else, the reason is why ad hoc algorithms work is that all of these techniques result in simple exponential filters for estimating the states. The actual filter that performs the state estimates is an exponential filter. Incorrect parameter estimate results in having a slightly non optimum filter link. The worst your estimate of the parameter, the more non optimum it is. That is its effect and so it is upsetting, but not tragic. It is serious, but not catastrophic.

P. Tavella (IEN, Italy): You spoke of the possibilities of using different states to optimize the things, maybe short, medium and long term stability. In the definition, you describe but do you use different equations with different ways, and it seems to me that in this way you define two or three different in different time scales. How can you obtain a unit time scale - do you make a sort of frequency lock of one on the other.

S. Stein: No! They are not different scales. All three equations can be satisfied simultaneously. It is as simple as that.

P. Tavella: That if you solve the three equation, you will obtain the redefined ensemble time scale.

S. Stein: No! Remember you have a weighted average of the phases of the clock that is constrained. You have a weighted average of the frequencies that are constrained and you have a weighted average of the frequency aging that are constrained and they may all be constrained and achieve a single solution. Each of the states is independently perturbed. So we have one solution and there is no funny steering to combine them.

P. Tavella: In the equation for phase continuity in the prediction of the phase that you also put frequency and (?) or also phase.

S. Stein: You use the frequency in there to make the forecast. In the basic time scale equation we say the weighted average of the time of the clock with respect to the time scale is equal to the weighted average of the forecasts. That is the equation you wrote in your paper two years ago. It is the same except that I use "a" and you use "p". The forecast depends on the frequency estimate. The forecast is the old phase estimate, plus the old frequency estimate, times the interval delta, plus one half delta squared, times the old frequency aging estimate. We constrain independently now the current frequency estimate to be the weighted average of the frequency forecast. This equation and the basic equation are all required simultaneously - the same solution. I have lots of degrees of freedom.

R. Clark (USNO): When you have subgroups of clocks that are exposed to environmental

effects and that type. There is a correlation between groups of clocks. How is this treated and is that looked for and sometimes that is just entirely not obvious - things like preventive maintenance that is done on the first of the month or third Tuesday of the month with clocks that are not even in the same area.

S. Stein: I think the answer to that question is at the present time nobody has a tentative process for that kind of information. We do not.